

# Material Recognition from Local Appearance in Global Context

Gabriel Schwartz Ko Nishino  
 Department of Computer Science, Drexel University  
 {gbs25, kon}@drexel.edu

## Abstract

Recognition of materials has proven to be a challenging problem due to the wide variation in appearance within and between categories. Many recent material recognition methods treat materials as yet another set of labels like objects. Materials are, however, fundamentally different from objects as they have no inherent shape or defined spatial extent. This makes local material recognition particularly hard. Global image context, such as where the material is or what object it makes up, can be crucial to recognizing the material. Existing methods, however, operate on an implicit fusion of materials and context by using large receptive fields as input (i.e., large image patches). Such an approach can only take advantage of limited context as it appears during training, and will be bounded by the combinations seen in the training data. We instead show that recognizing materials purely from their local appearance and integrating separately recognized global contextual cues including objects and places leads to superior dense, per-pixel, material recognition. We achieve this by training a fully-convolutional material recognition network end-to-end with only material category supervision. We integrate object and place estimates to this network from independent CNNs. This approach avoids the necessity of preparing an infeasible amount of training data that covers the product space of materials, objects, and scenes, while fully leveraging contextual cues for dense material recognition. Experimental results validate the effectiveness of our approach and show that our method outperforms past methods that build on inseparable material and contextual information.

## 1. Introduction

Material recognition is an inherently challenging problem, primarily due to the large variation in appearance between different instances of a given material and between different materials. There has been significant recent progress in terms of accuracy on benchmark datasets. Most methods proposed to achieve such results, however, essentially treat material recognition as object recognition with




Figure 1. Material recognition requires accurate sources of both local and global information. Even if a material is not uniquely determined by local appearance, local cues can constrain the choice of possible materials. Here, the local appearance suggests a smooth white material such as ceramic or plastic, and the global appearance is that of a mug. Rather than mix these two sources of information implicitly by just using large patches as input, we separate and then combine local materials and global object and scene context to improve material recognition.

different categories. They often use large image patches that cover parts or whole objects and scenes, as one would when performing object recognition, which inevitably mix visual cues of materials and other image context, mainly of objects.

Material recognition is fundamentally different from object recognition. Adelson [1] alludes to this difference in his discussion of “things” vs. “stuff”. While materials underlie both “things” and “stuff,” the key difference between them highlights the critical difference between objects and materials. Similar to “stuff”, and unlike “things” (objects), materials may not necessarily be recognized by having a particular shape. A cup, an object with a typical cylindrical shape, is often made of ceramic, a material. The fact that an object is a cup can be used as a cue to recognize the material as ceramic, and likewise the presence of ceramic may suggest that an object might be a cup. “Cup” and “ceramic” are not, however, interchangeable. Not all ceramic “things” are cups and relying on shape cues to recognize ceramic, which

past methods inevitably do, is a fundamentally limited approach. Only if we can recognize ceramic without realizing that the object is a cup, can we combine the fact it is made of ceramic with a cylindrical shape to help recognize that it might be a cup.

For this, materials need to be recognized locally (i.e., without seeing the object it makes up or the scene in which it lies), and other “global” context including objects and scenes recognized separately can help eliminate remaining uncertainty arising from strictly local appearance cues. Recognition of materials from local appearance cues becomes even more essential when recognizing “stuff” such as towels, water, and bushes that do not have canonical shapes.

Mixing context and material categories during material recognition implicitly relies on the underlying recognition framework to disentangle these concepts. Bell *et al.* [2] come to the conclusion that, “Training on a dataset which includes the surrounding context is crucial for real-world material classification,” but when simply given large patches as input and materials as output, we may only speculate as to the actual importance of context.

On the other hand, there have been recent attempts to study the separation of materials from other image context. Hu *et al.* [6] briefly investigated the correlations between objects and materials. Schwartz and Nishino [13, 14] proposed to recognize materials from small local image patches taken from inside the boundary of object regions and do not contain any shape cues to achieve per-pixel recognition based on visual material traits (e.g., “smooth,” “hard,” and “organic”) as discriminative internal representations that transcend material categories. Although these methods successfully demonstrate the power of local material recognition, they do not show the integration of rich image context for combined reasoning of materials in images.

Global image context, including both what the object is and where it is, can provide rich information to narrow down what things and stuff are made of. A city street, for example, is likely to contain asphalt, rubber, glass, and metal materials. Figure 2 shows examples of actual correlations between materials and context given ground-truth objects from the MS COCO database [10], along with global context in the form of place category predictions from the MIT Places CNN [21]. The challenge in exploiting these context cues is that if we were to simply attempt to obtain exhaustive annotations for materials, objects and places, we would be searching a product space with an extremely large number of combinations.

In this paper, we introduce a material recognition framework that fully leverages local appearance and global contextual cues to recognize materials at each pixel in an image. By relying on accurate global context that may be extracted from any image, we avoid having to obtain annotations for

$\{Places\} \times \{Objects\} \times \{Materials\}$ . Materials are an inherently local property, and as such we aim to produce dense per-pixel material category predictions. We separate materials from objects and other context by training a full-resolution dense-output material CNN on small local image patches. We then introduce explicit accurate context cues, in the form of object and place category predictions, to higher levels of the network. This allows us to provide accurate context rather than the implicit and uncertain context present in large patches, and also avoids the tradeoff between spatial resolution and context observed in [2].

We group the context categories into a logical hierarchy and investigate the effect of hierarchical context level on material recognition. We find that each additional form of context we introduce provides an independent increase in material recognition accuracy, and that finer-grained context is better for material recognition. We also investigate the ideal level, in terms of the hierarchical levels of the CNN recognition framework, at which to introduce context. Intuitively, objects and places are high-level concepts. Our results agree with this, as we find that context is best used when it we introduce it at the highest level of the network.

Our results show that the explicit separation and (re-)integration of image context significantly improves local material recognition accuracy. We quantitatively evaluate the accuracy of our method and find that it outperforms previous approaches that implicitly mix materials with object and place context. The human visual system is inherently a combination of top-down and bottom-up processes [9, 11]. Our approach may be considered a first step towards a full integration of bottom-up (materials) and top-down (context) information for material recognition. Our paper focuses on integrating sources of top-down context and the results suggest that a full integration, where the goal is to recognize both materials and context, would be a promising next step for future work.

## 2. Related Work

Material recognition is usually done at an image patch level. These patches are, however, relatively large in most existing works: they span a significant area of the scene, sometimes even the entire image, covering parts of or entire objects. Sharan *et al.* [16] introduced the earliest form of such classification with the Flickr Materials Database (FMD). In the FMD, the image patch is the entire image, and each image contains a single primary material of interest, similar to image classification. Recently, Bell *et al.* [2] demonstrated per-pixel material classification using a large-scale annotated training data, the Materials in Context (MINC) dataset, and a combination of CNN and CRF models for classification. Their method uses a large image patch for each pixel, roughly a quarter of the entire image, which inevitably mix in object or place context to material appear-

ance. Wang *et al.* [18] also demonstrate accurate dense per-pixel material predictions using 4D light field images. Zhang *et al.* [20] have recently shown impressive performance on the FMD, but their results focus only on single patch predictions. These methods mix materials and context interchangeably throughout the recognition pipeline, when they would be better-used in a factorized form as we show in this paper.

Dense prediction, outputting a value or category prediction for each pixel, has been extensively studied in the context of object recognition and object semantic segmentation. Object recognition datasets, such as ImageNet [12] or MS COCO [10], often contain many (80-10,000) categories. Despite this, state-of-the-art semantic segmentation methods such as DeepLab [4] focus on only a small subset of coarse-grained categories. A notable and relevant exception is the recent ADE20k dataset, scene parsing challenge, and associated models [22]. The dataset contains many fully-segmented images, and the challenge defines a set of 150 categories for semantic segmentation. We are not merely performing semantic segmentation, though we do aim to produce dense material predictions. Our goal is to properly recognize materials from local appearance. We do, however, find the ADE20k models to be ideal sources of per-pixel object category context information.

The use of context as a means to reduce ambiguity, whether in materials or other cases, appears promising. Hu *et al.* [6] showed that a simple addition of object category predictions as features could potentially improve material recognition. On an unrelated topic, Iizuka *et al.* [7] use scene place category predictions to improve the accuracy of grayscale image colorization. Our work, in contrast to these previous methods, takes advantage of multiple sources of context and investigates the ideal granularity of context categories. Within the framework of a Convolutional Neural Network (CNN), we evaluate how the hierarchical level at which we introduce context influences the accuracy of the corresponding material predictions.

### 3. Local Material Recognition

We aim to leverage scene context, such as objects and places, to improve dense per-pixel material recognition. Our first step is to ensure that, when recognizing materials, we are in fact dealing with just materials and not an implicit fusion of materials and context. Schwartz and Nishino [13, 14] have proposed to achieve such a separation by recognizing materials using only small local image patches inside the boundary of objects as input, thereby avoiding any influence from context derived from object shape or other global features. They have also recently introduced a dataset aimed at local material recognition, with carefully-selected categories chosen from a material hierarchy [15], and material annotations that respect object

boundaries. Although their model is not ideal for dense per-pixel material recognition, we do take advantage of their material dataset and build upon the concept of local recognition from small patches.

Our goal is to integrate materials with context that may be partially global (objects with large spatial extent but defined boundaries) or fully global (scene place categories, one per image). The frameworks introduced in [13, 14] are only able to make dense pixel-wise predictions in a sliding window fashion, and do not offer any logical point at which to introduce global context. To address this, we build a fully-convolutional CNN architecture, based on the VGG-16 network of Simonyan and Zisserman [17] with modifications to enable us to output dense full-resolution material predictions with integrated global context. Bell *et al.* [2] have previously investigated a similar architecture for material recognition. They, however, rely on large (24% image size) training patches and a loosely-defined set of material categories (e.g., carpet as a material) that collectively mix up materials and objects. In contrast, all of our training is done with small local material image patches taken inside object boundaries. Section 5.1 describes the model architecture.

## 4. Conditional Distributions of Materials and Context

We have an intuitive understanding that if one knows an object is, for example, made of metal, then it may be a knife or a car, but probably not a piece of clothing. Likewise, if we know an object is a cup, then it is likely made of glass, plastic, or ceramic. We can quantitatively evaluate the informative nature of context, such as object and place categories, by estimating the conditional probability distributions of materials given each possible category of context. If our intuition is correct, then these distributions should be very discriminative (have a low entropy relative to the corresponding discrete uniform distribution).

### 4.1. Object Context

We can get an initial idea as to how discriminative context is by using ground-truth object masks and corresponding materials to estimate the conditional distribution  $p(M|O)$ , where  $M$  is the material category and  $O$  is the object category. We use the material database of [15] as they include images from databases that contain object category map annotations (particularly, MS COCO [10]). To compute the conditional probabilities, we take each image with material annotations and find the object exhibiting each material as indicated by the COCO ground truth. Figure 2 shows conditional material probabilities  $p(M|O = o)$  for a few selected object categories  $o$ . The entropy for the discrete uniform distribution over 16 categories is 2.77,




Figure 2. The conditional distributions of materials given ground-truth object categories (top row) and predicted places (bottom row) are highly discriminative. Many context categories exhibit only a small set of materials. Some outliers are inevitable as the ground-truth COCO segmentation masks do not perfectly conform to actual object boundaries in the image. Places do not offer the same very strong divisions of materials as objects do. Their distributions are, however, still valuable as shown both by their entropy and the resulting material recognition accuracy based on place context.

and we can see the entropy given true object categories is significantly lower.

## 4.2. Place Context

Objects are defined somewhat locally (at the level of groups of pixels, but still globally compared to the local material appearance we model) and tend to exhibit only a small set of materials. In contrast, places are single scene-wide properties and can encompass many objects and materials. Despite this, we expect that places can still provide useful cues to disambiguate local materials. Ceramic and paper, for example, are often both flat white surfaces. Without seeing a specular highlight, it may be difficult to distinguish the two given only a small local patch. If, however, we know that the image patch originates from an image of a classroom, it is much more likely that the patch contains paper.

In the absence of ground-truth place category annotations, we rely on predictions from the MIT Places CNN [21]. Figure 2 contains examples of the conditional distributions  $p(M|P = p)$  for a few choices of place category  $p$ . While they are not uniformly as discriminative as object categories, they still do provide some useful cues. Botanical gardens, for example, tend to contain plants as




Figure 3. Distribution of ratios of predicted vs. true material probabilities on a set of incorrect predictions. The median of the ratios is 3.8 and many are higher, indicating that the output of the softmax layer will not easily be affected by simple multiplication.

one would expect, and images of crosswalks contain asphalt, metal, and rubber (roads, cars).

At least in the case of objects, and perhaps places as well, the conditional distributions are so discriminative as to suggest that we might simply multiply these distributions with the predictions of an existing material recognition model and achieve improved accuracy. We initially investigated this applied to a material recognition CNN as a baseline for comparisons. We, however, found that simple multiplication made a negligible difference in the accuracy. This is due to the fact that many of the mispredictions we might hope to correct are too strongly-predicted for simple multiplication to have any effect, as seen in Figure 3.

## 5. Integrating Global Context and Local Material Appearance

While the context cues from objects and places are highly discriminative, we cannot simply treat them as a prior on material occurrence and multiply them with a model’s prediction. The model must instead have the context available during training, so that the context may influence the material predictions. Such an observation is consistent with the general idea of leveraging top-down feedback with bottom-up recognition, for instance, as demonstrated with object detection [5] and human pose estimation Carreira *et al.* [3]. Here, we are obtaining the top-down information in the form of object and place context. We treat the set of predicted context category probabilities, obtained from state-of-the-art networks for scene recognition and object semantic segmentation, as an additional feature in a dense per-pixel material CNN. By concatenating these probabilities with the high-level features in the network prior to output, we may take full advantage of the strong



material recognition cues available in global image context.

### 5.1. Context Integration Network

As shown in Figure 4, our model is based on the VGG-16 architecture of Simonyan and Zisserman [17] with a few key modifications to enable dense prediction from local material image patches with added global image context. To enable dense prediction, we train a fully-convolutional form of the network where all fully-connected layers are replaced with convolutions. We also add a series of upsampling layers to match the input and output resolutions.

We find that the level of downsampling in the original VGG-16 network is not compatible with local material recognition. The minimum patch size is constrained by the downsampling factor, and we would like to use small<sup>1</sup> image patches to maximize the separation between local materials and global context. To avoid too much downsampling, we remove the last set of pooling and filtering from the network. In order to train on datasets where densely-segmented material ground truth may not always be available, we compute the softmax loss function at each pixel. The loss function is only evaluated at pixels where there is a known material. In this way we are able to take advantage of segmented material images without requiring completely dense annotation.

The dense predictions from the network arise from a set of upsampling layers. Each upsampling layer performs a convolution with strided output to double the output resolution. Kim *et al.* [8] proposed the use of recursive convolution, repeated convolution with shared filter weights, for single-image super-resolution. The upsampling at the end of a dense prediction network may be viewed as a form of super-resolution, and we find that the use of recursive convolutions with a small recursion depth (3 in our case) after each upsampling step provides a small but consistent (2%) increase in accuracy.

For global image context integration, we treat estimated context category probabilities as additional features and concatenate them with existing features in the network. For places, we only have a single set of category probabilities for the entire image. We replicate these values across the image to match the image dimensions at the level where the context is introduced. This is similar to the colorization work of Iizuka *et al.* [7]. We, however, introduce per-pixel context in the form of object semantic segmentation and find that the ideal layer for introduction of the context is in fact higher in the network (see Section 5.4 below).

### 5.2. Extracting Per-Pixel Global Context

To fully leverage the separation of and later integration of materials and context, we must obtain the context from accurate sources. Here, we focus on context from per-pixel

| Hierarchy Level | Accuracy (% , per-pixel avg.) | Entropy |
|-----------------|-------------------------------|---------|
| High Level      | 51.0                          | 2.51    |
| Mid Level       | 53.3                          | 2.40    |
| Low Level       | 54.5                          | 2.27    |
| All Places      | 57.2                          | 1.91    |

Table 1. Place categories have an associated hierarchy that are grouped by various attributes such as indoor/outdoor or man-made/natural. We might expect that the grouped places at each level of the hierarchy contain more or less discriminative combinations of materials. Fine-grained places may not appear in many images, but coarse grained categories may offer little in the way of material recognition cues. We in fact find that despite this, the finest category granularity offers the best material recognition performance. In this case, the 205 place categories are both fine-grained and sufficiently well-distributed across training examples.

object category predictions (semantic segmentation) and single scene-wide place category predictions. For places, we rely on the MIT Places CNN [21]. A key feature of their places database is that the categories are hierarchically organized. As we will show below when we investigate the effects of place category granularity, the number of context categories is important: fine-grained context provides more useful information than simple, higher-level context groups. Most existing object semantic segmentation methods train their models on only a relatively small set of high-level object categories. A notable exception is the ADE20k dataset [22]. The dataset contains over 2,000 object categories. The MIT Scene Parsing Challenge, which relies on this dataset, selects 150 categories for semantic segmentation and it is these categories we use as our per-pixel object context.

### 5.3. Hierarchical Place Context

As part of their SUN database for scene and object recognition, Xiao *et al.* [19] define a hierarchy of place categories. This hierarchy raises the question of whether any particular context granularity is more or less useful for material recognition. On one hand, having an extremely fine set of place categories might mean that few training examples would appear from certain places. At the other extreme, the coarsest division of places could only provide very general cues as to which materials may be present.

To evaluate the importance of place granularity, we compute material recognition accuracy scores using only place context at each level of the SUN places hierarchy. We adapt their hierarchy to the place categories recognized by the MIT places CNN and treat nodes within each level of the hierarchy as place categories. The highest level is the simple division of indoors vs. outdoors, mid-level categories deal with distinctions such as commercial and residential buildings, or mountains and forests, and the lowest level includes smaller groups such as entertainment or religious places. Results in Table 1 show that accuracy

<sup>1</sup>48 × 48

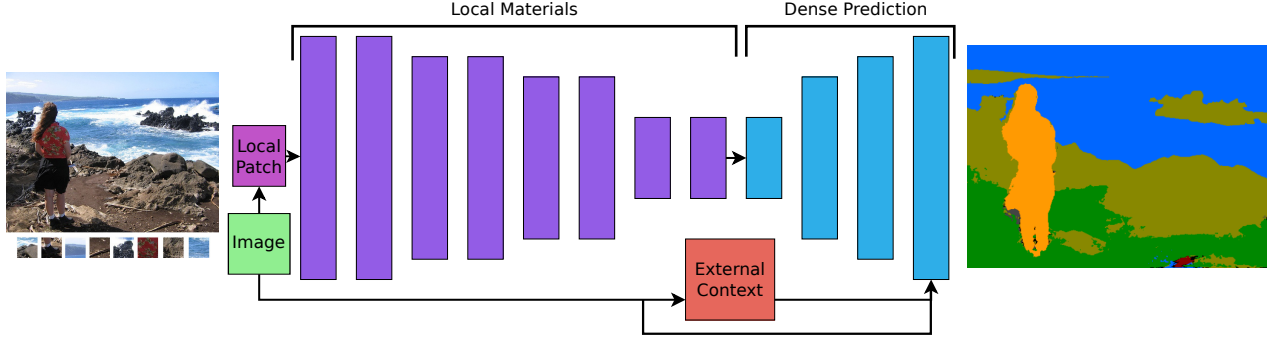


Figure 4. We integrate local materials with scene context in a CNN framework based on the VGG-16 architecture. We employ fewer pooling steps to allow us to train on small local image patches, and we add upsampling layers with recurrent convolutions (blue) inspired by single-image super-resolution. Per-pixel global context, in the form of place and object category probabilities predicted from the input image via existing models, is treated as an additional feature and concatenated with the upsampled result prior to final prediction. For non-localized context like place categories, we broadcast the probabilities across the entire image. For improved spatial resolution, we also add a skip connection between the input image and final output.

increases with place category granularity: more detailed place categories provide more discriminative information for material recognition. Computing the entropy of the conditional distributions  $p(M|P_i)$  for place category set  $P_i$  at hierarchy level  $i$  supports these results.

#### 5.4. Context Integration Level

In the case of a full top-down/bottom-up integration of materials and context, in which the goal was to predict both material and context categories from an image, we would expect that the context should be introduced at the lower levels of the CNN hierarchy to allow information from two categories to influence each other. In our case, however, we are focusing on combining global, presumed-accurate, context with material predictions. This allows us to avoid the requirement of place and object annotations for each material. The product space of objects, places, and materials is large, and obtaining such exhaustive combinations of annotations would be infeasible. It is the case, however, that we thus do not have annotations for input to a place/object category loss function. There is a possibility that the ideal level for context introduction is now not the lowest, and we investigate that possibility by comparing per-pixel average material recognition accuracies when introducing the context at varying levels of the network.

As we have seen given the preliminary investigation of context above, objects and places provide strong cues for the presence or absence of materials in a scene. If introduced in the wrong way, however, these cues may cause overfitting, resulting in predictions that rely solely on context. As the graph in Figure 5 shows, the ideal level for context introduction is in fact high in the network hierarchy, immediately before final classification. As the context is already extremely discriminative (as we have shown above), it should require little in the way of additional features in order to fully exploit it. Furthermore, in the case of object context, much of the spatial information present in the ob-




Figure 5. We have shown that objects and places can offer highly discriminative cues for material recognition. The integration, however, introduces a potential for overfitting by relying too heavily on the context. By plotting material recognition accuracy at various context introduction levels, we see that the best level for context introduction is the penultimate layer in the network. Object context also contains spatial information which is lost when downsampled to match the size of features lower in the network. This is in contrast to the findings of [7], where their context had no spatial information and was introduced prior to upsampling.

ject category maps is lost when the context is downsampled to be compatible with lower levels of the network.

## 6. Dense Material Recognition with Global Context

Table 2 contains a breakdown of the contributions for each form of global context. Most importantly, the contributions from objects and places are similar and the combination of the two outperforms either individual context source. This suggests that the objects and places are providing orthogonal sources of information and are both critical to

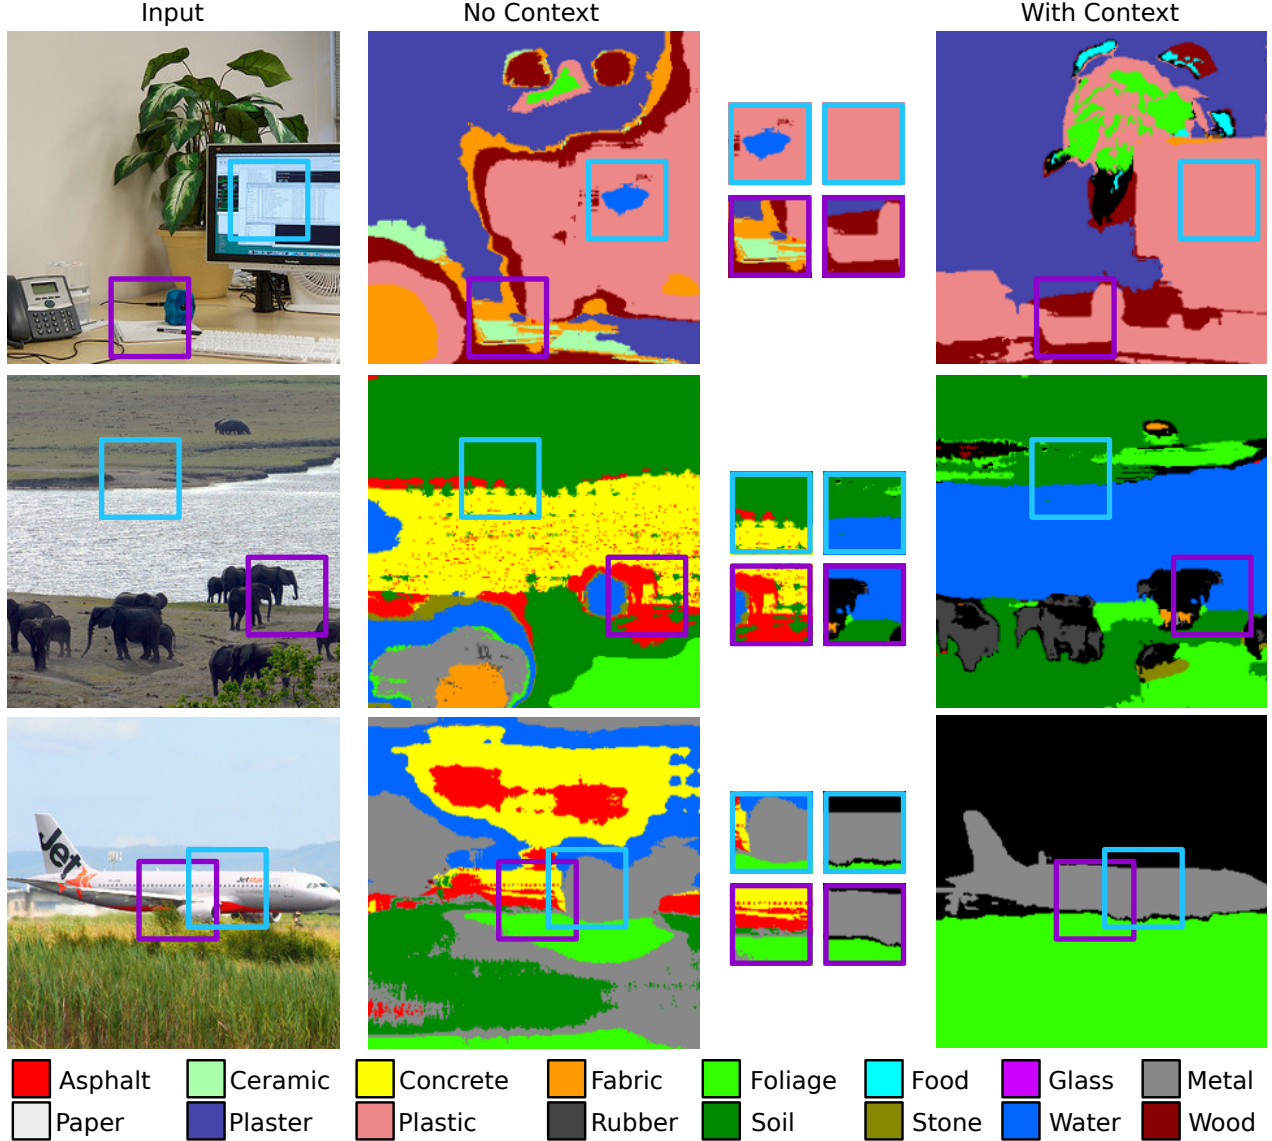


Figure 6. These examples show that context helps disambiguate materials when local information is not sufficient. In the first set of insets, the white paper and plastic could be mistaken for ceramic, but context suggests otherwise. In the second, the water has a local appearance similar to concrete. Global context suggests that this is unlikely. In the third set, we see that the airplane body is incorrectly recognized as concrete and asphalt due to the lack of characteristic specular highlight. Again, context fixes this error. Black pixels indicate regions where the material is uncertain (low probability), often when the pixels do not contain a known material.

the accurate recognition of materials. A cup, for example, may be made of glass or plastic. If, however, you are in a bar, then it is much more likely to be made of glass.

We argued that methods relying on large image patches can only take advantage of context in an implicit fashion, by using whatever object shape cues and scene features happen to be visible in the large image patch. In contrast, our proposed approach explicitly introduces known-discriminative forms of context to a network for purely local material recognition in an explicit fashion. Results in Table 3 show that our approach outperforms the large-patch-based (VGG-16 + CRF) method of Bell *et al.* [2] when evaluating per-

pixel average material accuracy on the overlapping categories of the local materials database of [15] and MINC [2]. The MINC model’s per-pixel accuracy suffers on the local materials database [15] compared to their own database. These results suggest that the database in [15] contains more diverse and challenging images: a large part of the MINC database comes from real-estate photographs [2] and thus are inevitably biased in materials.

We can readily see in Figure 6 that the context helps disambiguate materials that may be difficult to recognize from only local information. When metal does not exhibit specular highlights or reflections, as is the case with the




Figure 7. Additional examples of dense material recognition with global context show that the spatial information in the object context improves the precision of the locally-recognized materials. Each row contains three input images in the leftmost column and corresponding outputs to the right. It is important to note that neither skin nor sky appear as materials in the definition of [15]. Skin is a unique case of a material that is visible only on one object category (people) in most databases, and the sky is not a material.

| Context          | Accuracy (% , per-pixel avg.) |
|------------------|-------------------------------|
| None             | 50.8                          |
| Only Places      | 57.2                          |
| Only Objects     | 56.1                          |
| Places + Objects | <b>64.9</b>                   |

Table 2. By comparing per-pixel average accuracy with no context, as well as with each separate form of additional context, we see that the two forms of context (objects and places) each have a strong effect on the recognition of materials. The combination of objects and places is also significantly higher than either of the two alone, suggesting that knowing both objects and places provides unique cues not present in either individual category group.

| Method                     | Accuracy (% , per-pixel avg.) |
|----------------------------|-------------------------------|
| MINC [2]                   | 60.0                          |
| Ours<br>(Places + Objects) | <b>65.4</b>                   |

Table 3. Our initial argument was that large-patch-based methods can only implicitly use any available context and thus cannot make the best use of said context. We compare our per-pixel average material accuracy on the material database of [15] with the high-performing VGG16-based approach of [2] to show that explicit addition of context outperforms the implicit approach. As in [2], we measure accuracy on overlapping categories only.

airplane body, the flat white surface offers little in the way of local recognition cues. Knowing either that the scene is an airport or that the current pixel belongs to a plane removes this ambiguity. Likewise, in the natural scene with elephants, the combination of high-frequency waves and

specular reflection causes the water to appear like concrete. Scene context makes it clear that concrete would be unlikely in this case. In general, the predictions are accurate subject to the limitations of the training data. Skin is not a material in the dataset of [15], and thus skin is often classified as the surrounding fabric. Sky is not a material and the predictions for sky are determined largely by context (ex. metal at airports, water over the ocean).

## 7. Conclusion

Our results show that we can successfully separate materials from their surrounding context, and combine those materials with highly-discriminative forms of global context. Such a combination outperforms previous methods which implicitly rely on context being available in a large input image patch. Additionally, we perform a detailed investigation into the ideal granularity for context in material recognition as well as the hierarchical level at which the context should be introduced into a CNN framework. Additional qualitative examples in Figure 7 show that the combination of local materials and global context results in accurate material predictions in the face of local ambiguity.

Based on the work of Mumford *et al.* [9, 11] and more recent results by Carreira *et al.* [3], we think that a full factorization and integration of top-down context with bottom-up materials is an interesting area for future work. The results we find in this paper are a promising first step towards such a goal.



**Acknowledgments** This work was supported by the Office of Naval Research grant N00014-16-1-2158 (N00014-14-1-0316) and the National Science Foundation awards IIS-1353235 and IIS-1421094.

## References

- [1] E. H. Adelson. On Seeing Stuff: The Perception of Materials by Humans and Machines. In *SPIE*, pages 1–12, 2001.
- [2] S. Bell, P. Upchurch, N. Snavely, and K. Bala. Material Recognition in the Wild with the Materials in Context Database. In *CVPR*, 2015.
- [3] J. Carreira, P. Agrawal, K. Fragkiadaki, and J. Malik. Human Pose Estimation with Iterative Error Feedback. volume abs/1507.06550, 2015.
- [4] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *arXiv*, abs/1606.00915, 2016.
- [5] B. Epshtein, I. Lifshitz, and S. Ullman. Image Interpretation by a Single Bottom-Up Top-Down Cycle. *PNAS*, 105(38), 2008.
- [6] D. Hu, L. Bo, and X. Ren. Toward Robust Material Recognition for Everyday Objects. In *BMVC*, pages 48.1–48.11, 2011.
- [7] S. Iizuka, E. Simo-Serra, and H. Ishikawa. Let there be Color!: Joint End-to-end Learning of Global and Local Image Priors for Automatic Image Colorization with Simultaneous Classification. In *SIGGRAPH*, volume 35, pages 110:1–110:11, 2016.
- [8] J. Kim, J. K. Lee, and K. M. Lee. Deeply-Recursive Convolutional Network for Image Super-Resolution. In *CVPR*, pages 1637–1645, 2016.
- [9] T. S. Lee and D. Mumford. Hierarchical Bayesian Inference in the Visual Cortex. *J. Opt. Soc. Am. A (JOSA A)*, 20(7):1434–1448, 2003.
- [10] T.-Y. Lin, M. Marie, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: Common Objects in Context. In *ECCV*, 2014.
- [11] D. Mumford. On the Computational Architecture of the Neocortex II: The Role of Cortico-Cortical Loops. *Biological Cybernetics*, 66:241–251, 1992.
- [12] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [13] G. Schwartz and K. Nishino. Visual Material Traits: Recognizing Per-Pixel Material Context. In *Color and Photometry in Computer Vision (Workshop held in conjunction with ICCV’13)*, 2013.
- [14] G. Schwartz and K. Nishino. Automatically Discovering Local Visual Material Attributes. In *CVPR*, pages 3565–3573, 2015.
- [15] G. Schwartz and K. Nishino. Discovering Perceptual Attributes in a Deep Local Material Recognition Network. *arXiv*, abs/1604.01345, 2016.
- [16] L. Sharan, R. Rosenholtz, and E. Adelson. Material Perception: What Can You See in a Brief Glance? *Journal of Vision*, 9(8):784, 2009.
- [17] K. Simonyan and A. Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *ICLR*, pages 1–14, 2015.
- [18] T.-C. Wang, J.-Y. Zhu, H. Ebi, M. Chandraker, A. A. Efros, and R. Ramamoorthi. A 4D Light-Field Dataset and CNN Architectures for Material Recognition. In *ECCV*, 2016.
- [19] J. Xiao, J. Hays, K. Ehinger, A. Oliva, and A. Torralba. SUN Database: Large-scale Scene Recognition from Abbey to Zoo. In *CVPR*, 2010.
- [20] Y. Zhang, M. Ozay, X. Liu, and T. Okatani. Integrating Deep Features for Material Recognition. *arXiv*, abs/1511.06522, 2015.
- [21] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning Deep Features for Scene Recognition using Places Database. In *NIPS*, 2014.
- [22] B. Z. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba. Semantic Understanding of Scenes through ADE20K Dataset. *arXiv*, abs/1608.05442, 2016.